

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

- 1 1. (Currently amended) A method of predicting cache performance,
2 comprising:
3 storing data references applied to an operational data cache in a data
4 processing environment;
5 applying said data references to a cache simulator configured to
6 simultaneously simulate a plurality of caches of different sizes, wherein said cache
7 simulator comprises multiple segments and each said simulated cache comprises
8 one or more of said segments; and
9 generating for each of said plurality of simulated caches an estimate of
10 performance based on said simulation, wherein said generating comprises:
11 for a first simulated cache in said plurality of simulated caches:
12 calculating the number of hits in all segments of
13 said cache simulator that do not comprise part of said first
14 simulated cache,
15 adding the number of said misses of all of said
16 segments to said calculated number of hits to produce an
17 initial estimated miss rate for said first simulated cache,
18 calculating a correction factor to apply to said
19 initial estimated miss rate for said first simulated cache, and
20 multiplying said initial estimated miss rate for said
21 first simulated cache by said correction factor to yield a
22 predicted miss rate for said first simulated cache;

23 wherein each application of one of said data references to said cache
24 simulator causes either a hit in one of said segments or a miss of every said
25 segment.

1 2. (Original) The method of claim 1, further comprising:
2 dynamically adjusting the size of said operational data cache to match the
3 size of one of said simulated caches.

1 3. (Original) The method of claim 1, wherein said data processing
2 environment is a database management system and said operational data cache is a
3 buffer cache configured to cache data in said database management system.

1 4. (Original) The method of claim 3, wherein said storing data references
2 comprises:
3 receiving a first data reference at said operational data cache;
4 storing said first data reference in a trace buffer for use in said simulation
5 of said plurality of caches if said trace buffer is not full; and
6 discarding said first data reference if said trace buffer is full.

1 5. (Original) The method of claim 1, wherein said cache simulator
2 comprises a list of simulated buffers, and each said simulated buffer is configured
3 to store:
4 an identifier of a data item; and
5 an identifier of said segment of said cache simulator in which said
6 simulated buffer is located.

1 6. (Original) The method of claim 5, wherein said applying comprises:
2 retrieving a first stored data reference;

3 identifying a first data item referenced in said first data reference; and
4 searching said cache simulator for a buffer in which an identifier of said
5 first data item is stored.

1 7. (Original) The method of claim 6, further comprising:
2 if a first buffer is found in said cache simulator that stores an identifier of
3 said first data item:
4 incrementing a hit counter for said segment in which said
5 first buffer is found;
6 moving said first buffer to a head of said cache simulator;
7 and
8 updating said stored segment identifiers of said buffers as
9 necessary.

1 8. (Original) The method of claim 6, further comprising:
2 if no buffer is found in said cache simulator that stores an identifier of said
3 first data item:
4 incrementing an absolute miss counter;
5 storing an identifier in a first buffer in said cache simulator;
6 moving said first buffer to a head of said cache simulator;
7 and
8 updating said stored segment identifiers of said buffers as
9 necessary.

1 9-10 (Canceled).

1 | 11. (Currently amended) The method of ~~claim 10~~claim 1, wherein said
2 correction factor comprises the ratio of misses incurred during application of said

3 data references to said operational data cache to an initial estimated miss rate for a
4 second simulated cache in said plurality of simulated caches;
5 wherein the number of buffers in said second simulated cache matches the
6 number of buffers in said operational data cache.

1 12 (Canceled).

1 13. (Original) A method of simulating the performance of multiple caches
2 of different sizes, comprising:
3 storing references to data items received during operation of a database
4 management system;
5 maintaining a multi-segmented cache simulator comprising simulated
6 buffers configured to store data identifiers, wherein each of said multiple caches
7 comprises a set of said segments different from the other caches;
8 for each of said stored references:
9 searching said cache simulator for a first simulated buffer
10 storing an identifier of said referenced data item;
11 if said first simulated buffer is found in said segmented
12 memory:
13 incrementing a hit counter for said segment in
14 which said first simulated buffer is located; and
15 moving said simulated cached reference to the head
16 of said segmented memory;
17 if no simulated buffer storing an identifier of said
18 referenced data item is found in said cache simulator:
19 incrementing a miss counter;
20 storing an identifier of said data item
21 in a second simulated buffer; and

22 storing said second simulated buffer
23 at a head of said cache simulator; and
24 generating an estimate of the performance of each
25 of said multiple caches.

1 14. (Original) The method of claim 13, wherein said estimated
2 performance comprises a number of misses, wherein said generating comprises:
3 for a first cache of said multiple caches, calculating the number of said
4 referenced data items for which identifiers were not found in the cache simulator
5 segments comprising said first cache.

1 15. (Original) The method of claim 14, further comprising:
2 calculating a predicted number of misses for said first cache by
3 multiplying said estimated performance for said first cache by a correction factor.

1 16. (Original) The method of claim 15, wherein:
2 said data item references are received at an operational cache of said
3 database management system and result in a known number of misses in said
4 operational cache;
5 a second cache of said multiple caches comprises a number of simulated
6 buffers equivalent to the number of buffers comprising said operational cache; and
7 said correction factor is equal to the ratio of said known number of misses
8 to said estimated performance for said second cache.

1 17. (Original) The method of claim 16, further comprising:
2 dynamically altering the size of said operational cache to the size of said
3 first cache, wherein said predicted number of misses for said first cache is less
4 than said known misses in said operational cache.

1 18. (Original) The method of claim 13, wherein said storing comprises:
2 receiving a first reference to a data item at said database management
3 system;
4 if a trace memory configured to store said first reference is full, discarding
5 said first reference; and
6 otherwise, storing said first reference in said trace memory without locking
7 said trace memory.

1 19. (Original) The method of claim 13, wherein said data item references
2 are received at an operational cache of said database management system, further
3 comprising:
4 identifying the number of operational cache misses resulting from
5 application of said data item references to said operational cache; and
6 dynamically altering the size of said operational cache to the size of one of
7 said multiple caches having a superior estimated performance.

1 20. (Original) A computer readable storage medium storing instructions
2 that, when executed by a computer, cause the computer to perform a method of
3 predicting cache performance, the method comprising:
4 storing data references applied to an operational data cache in a data
5 processing environment;
6 applying said data references to a cache simulator configured to
7 simultaneously simulate a plurality of caches of different sizes, wherein said cache
8 simulator comprises multiple segments and each said simulated cache comprises
9 one or more of said segments; and
10 generating for each of said plurality of simulated caches an estimate of
11 performance based on said simulation, wherein said generating comprises:
12 for a first simulated cache in said plurality of simulated caches:

13 calculating the number of hits in all segments of said cache
14 simulator that do not comprise part of said first simulated cache,
15 adding the number of said misses of all of said segments to
16 said calculated number of hits to produce an initial estimated miss
17 rate for said first simulated cache,
18 calculating a correction factor to apply to said initial
19 estimated miss rate for said first simulated cache, and
20 multiplying said initial estimated miss rate for said first
21 simulated cache by said correction factor to yield a predicted miss
22 rate for said first simulated cache;
23 wherein each application of one of said data references to said cache
24 simulator causes either a hit in one of said segments or a miss of every said
25 segment.

1 21-27 (Canceled).